

B. REMARKS

The Examiner is thanked for the performance of a thorough search. By this amendment, Claims 2, 17 and 32 have been canceled. Hence, Claims 1, 3-16, 18-31 and 33-40 are pending in this application. The amendments to the claims do not add any new matter to this application. All issues raised in the Office Action mailed October 11, 2006 are addressed hereinafter.

REJECTION OF CLAIMS 1-20 UNDER 35 U.S.C. § 101

Claims 1-20 are rejected under 35 U.S.C. § 101 as being directed to non-statutory subject matter. The stated basis for this rejection is that the method of Claim 1, the computer-readable medium of Claim 16 and the apparatus of Claim 31 do not produce any useful and tangible result. The rejection as set forth in the Office Action specifies that this rejection applies to Claims 1-20 and yet the accompanying remarks include remarks for apparatus Claim 31. These remarks therefore assume that the rejection was intended to apply to all pending claims, i.e., Claims 1-40, and not just Claims 1-20. This rejection is moot with respect to canceled Claims 2, 17 and 32.

It is generally accepted that most, if not all, database systems will, at some point in time, fail or “crash” as it is sometimes referred to. As a result, database systems conventionally include a mechanism for recovering from a crash. The recovery mechanism typically includes generating redo and undo records that are used to return the database to a transaction consistent state that existed just prior to the crash. In a transaction consistent state, the database reflects all the changes made by transactions that have committed and none of the changes made by transactions that have not committed. The redo records are used to reapplying all changes made to the database up to a point just prior to the crash. This process is conventionally referred to as “rolling forward.” These changes include all committed changes as well as any uncommitted changes that were recorded in the redo records prior to the crash. Then the undo records are used to remove uncommitted changes from the database, i.e., to remove changes made by transactions that were not committed at the time of the crash. This process is conventionally referred to as “rolling back” the database.

One significant issue with conventional database recovery mechanisms is that they are generally very computationally expensive to operate and can require a long time to complete

recovery, particularly when a large number of uncommitted transactions were active at the time of the crash. At least part of this is due to the use of a single recovery process to process the redo and undo records serially, which can require a significant amount of time when many changes need to be made. During this time, the database cannot be used for normal operations.

The approach for performing database recovery after a crash of an instance of a database as recited in method Claim 1, computer-readable medium Claim 16 and system Claim 31, respectively, reduces the amount of time required to perform database recovery. The time savings is achieved by using statistical data about dead, i.e., uncommitted, transactions to determine the number of recovery servers that should be used to recover the dead transactions. For example, in situations where a large number of undo blocks need to be recovered, multiple servers may be executed in parallel to perform the recovery to reduce the amount of time required to perform the recovery.

It is therefore respectfully submitted that the approach for performing database recovery after the crash of an instance of a database as recited in Claims 1, 16 and 31 provides a useful result because the approach can provide significant reductions in the amount of time required to recover a database after a crash and therefore make the database available sooner. This is particularly useful in situations where database “downtime” needs to be limited, for example in financial applications including banking and financial markets. It is also respectfully submitted that the approach for performing database recovery after the crash of an instance of a database as recited in Claims 1, 16 and 31 provides a tangible result because the result of the approach is a reduction of time required to perform database recovery, which can be measured and quantified, and has a direct effect on end users. It also allows users to more quickly resume using a database than they would have been able to using conventional recovery techniques.

In view of the foregoing, it is respectfully submitted that the approach for performing database recovery after the crash of an instance of a database as recited in Claims 1, 16 and 31 provides both a useful and tangible result and therefore satisfies the requirements of 35 U.S.C. § 101. Accordingly, reconsideration and withdrawal of the rejection of Claims 1, 3-16, 18-31 and 33-40 under 35 U.S.C. § 101 is respectfully requested.

REJECTION OF CLAIMS 1-40 AND 15 UNDER 35 U.S.C. § 103(a)

Claims 1-40 are rejected under 35 U.S.C. § 103(a) as being unpatentable over “ARIES-RRH: Restricted Repeating of History in the ARIES Transaction Recovery Method – Pages 718-727, 1991” by *Mohan* (hereinafter “*Mohan*”) in view of *Tada et al.*, U.S. Patent No. 5,544, 359 (hereinafter “*Tada*”). This rejection is moot with respect to canceled Claims 2, 17 and 32. It is respectfully submitted that Claims 1, 3-16, 18-31 and 33-40, as amended, are patentable over *Mohan* and *Tada*, considered alone or in combination, for at least the reasons provided hereinafter.

CLAIM 1

Claim 1, as amended, is directed to a method for performing database recovery after a crash of an instance of a database, wherein multiple transactions were active when the instance crashed. Claim 1 recites:

“identifying a plurality of dead transactions;
determining statistical data about the plurality of dead transactions;
determining that a particular number of recovery servers should be used to recover the plurality of dead transactions based on the statistical data; and
recovering the plurality of dead transactions using the particular number of recovery servers by executing the particular number of recovery servers in parallel.”

As previously mentioned herein, the approach for performing database recovery after a crash of an instance of a database as recited in method Claim 1 reduces the amount of time required to perform database recovery. The time savings is achieved by using statistical data about dead, i.e., uncommitted, transactions to determine the number of recovery servers that should be used to recover the dead transactions. The recovery servers are then executed in parallel to recovery the dead transactions.

It is respectfully submitted that Claim 1 recites one or more limitations that are not taught or suggested by *Mohan* and *Tada*, considered alone or in combination. For example, it is respectfully submitted that the Claim 1 limitations “determining statistical data about the plurality of dead transactions,” “determining that a particular number of recovery servers should be used to recover the plurality of dead transactions based on the statistical data” and “recovering the plurality of dead transactions using the particular number of recovery servers” are not taught or suggested by *Mohan* or *Tada*.

The Office Action asserts that these Claim 1 limitations are not taught or suggested by *Mohan* and are instead taught by *Tada* and the Applicant agrees that these limitations are not taught or suggested by *Mohan*. The Office Action refers to the text at Col. 13, lines 31-45, the text at Col. 13, line 65 through Col. 14, line 2 and the text at Col. 18, lines 35-45 in support of the assertion that the aforementioned limitations are taught by *Tada*. *Tada* discloses an approach for managing and processing log data by classifying the log data to aid in database recovery. The text at Col. 13, lines 31-45, the text at Col. 13, line 65 through Col. 14, line 2 describes how recovery operations are performed on a database to bring the database back to a point prior to a failure, so that operations can be resumed. This text also describes receiving and storing statistical information for system tuning. The text at Col. 18, lines 35-45 describes services that can be provided by classifying the log data in a certain manner. Some of the examples described include acquiring log data on a database-by-database manner and recovering log data that covers a shorter time span.

Applicant has studied these and other portions of *Tada* and does not find any teaching or suggestion of determining the number of recovery servers that should be used to recover a plurality of dead transactions based upon statistical data about the dead transactions. There is no mention or suggestion of using multiple servers to perform recovery or how the number of servers used to perform recovery might be determined. It is therefore respectfully submitted that the Claim 1 limitations “determining statistical data about the plurality of dead transactions,” “determining that a particular number of recovery servers should be used to recover the plurality of dead transactions based on the statistical data” and “recovering the plurality of dead transactions using the particular number of recovery servers” are not taught or suggested by *Tada*.

It is further respectfully submitted that the Claim 1 limitation “by executing the particular number of recovery servers in parallel” is also not taught or suggested by *Mohan* and *Tada*. This limitation was added to Claim 1 from Claim 2 in this reply. This limitation relates to how recovery is performed and more specifically, that dead transactions are recovered using the particular number of recovery servers by executing the particular number of recovery servers in parallel. Recovering dead transactions by executing recovery servers in parallel reduces the amount of time required to recovery the dead transactions and makes the database available sooner.

The Office Action asserts that this limitation is taught by *Mohan* at Page 722, Col. 1, first paragraph, lines 7-9 and at Page 720, Col. 1, second paragraph, lines 8-12. The text in *Mohan* at Page 722, Col. 1, first paragraph, lines 7-9 states “[t]hen, the undo of the loser transactions may be performed in parallel with the processing of new transactions.” After studying this text and other text in *Mohan*, it is respectfully submitted that this text does not stand for the proposition that the undo of loser transactions is performed using multiple recovery servers executing in parallel. Rather, this text teaches that the undo of loser transactions may be performed in parallel with the processing of new transactions. For example, a single recovery server may be used to perform the undo of loser transactions in parallel with a single server being used to process a new transaction.

The text in *Mohan* at Page 720, Col. 1, second paragraph, lines 8-12 describes that during the undo pass, that performing undos is not a conditional operation. The text further describes that this means that the ARIES system does not compare the page LSN of the affected page to the LSN of the log record to decide whether or not to undo the update. It is not clear why this portion of *Mohan* was relied upon with respect to this limitation of Claim 1 because there is nothing in this portion of *Mohan* related to recovering a plurality of dead transactions by executing recovery servers in parallel. There is also no suggestion in *Tada* of recovering a plurality of dead transactions by executing recovery servers in parallel and it is noted that *Tada* was not relied upon for teaching this limitation. In view of the foregoing, it is respectfully submitted that the Claim 1 limitation “by executing the particular number of recovery servers in parallel” is also not taught or suggested by *Mohan* and *Tada*.

In view of the foregoing, it is respectfully submitted that Claim 1 recites one or more limitations that are not taught or suggested by *Mohan* and *Tada*, considered alone or in combination, and that Claim 1 is therefore patentable over *Mohan* and *Tada*.

CLAIMS 3-15

Claims 3-15 all depend from Claim 1 and include all of the limitations of Claim 1. It is therefore respectfully submitted that Claims 3-15 are patentable over *Mohan* and *Tada* for at least the reasons set forth herein with respect to Claim 1. Furthermore, it is respectfully submitted that Claims 3-15 recite additional limitations that independently render them patentable over *Mohan* and *Tada*.

CLAIMS 16 AND 18-30

Claims 16 and 18-30 recite limitations similar to Claims 1 and 3-15, except in the context of computer-readable media. It is therefore respectfully submitted that Claims 16 and 18-30 are patentable over *Mohan* and *Tada* for at least the reasons set forth herein with respect to Claims 1 and 3-15.

CLAIMS 31 AND 33-40

Claims 31 and 33-40 recite limitations similar to Claims 1 and 3-15, except in the context of apparatuses. It is therefore respectfully submitted that Claims 31 and 33-40 are patentable over *Mohan* and *Tada* for at least the reasons set forth herein with respect to Claims 1 and 3-15.

In view of the foregoing, it is respectfully submitted that Claims 1, 3-16, 18-31 and 33-40 are patentable over *Mohan* and *Tada*. Accordingly, reconsideration and withdrawal of the rejection of Claims 1, 3-16, 18-31 and 33-40 under 35 U.S.C. § 103(a) as being unpatentable over *Mohan* in view of *Tada* is respectfully requested.

CONCLUSION

It is respectfully submitted that all of the pending claims are in condition for allowance and the issuance of a notice of allowance is respectfully requested. If there are any additional charges, please charge them to Deposit Account No. 50-1302.

The Examiner is invited to contact the undersigned by telephone if the Examiner believes that such contact would be helpful in furthering the prosecution of this application.

Respectfully submitted,

HICKMAN PALERMO TRUONG & BECKER LLP



Edward A. Becker

Reg. No. 37,777

Date: January 11, 2007

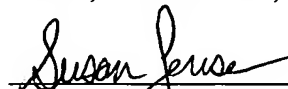
2055 Gateway Place, Suite 550
San Jose, CA 95110
Telephone: (408) 414-1204
Facsimile: (408) 414-1076

CERTIFICATE OF MAILING

I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail in an envelope addressed to: **Mail Stop Amendment**, Commissioner for Patents, P. O. Box 1450, Alexandria, VA 22313-1450

on January 11, 2007

by


Susan Jensen